

NPS-53-86-011

NAVAL POSTGRADUATE SCHOOL

Monterey, California



THE PERFORMANCE OF THE NEC SX-2 SUPERCOMPUTER
SYSTEM COMPATED WITH THAT OF THE
CRAY X-MP/4 AND FUJITSU VP-200

by

Raul H. Mendez

September 1986

Technical Report For Period

April - August 1986

Approved for public release; distribution unlimited

Prepared for: Naval Postgraduate School
Monterey, CA 93943-5000

FedDocs
D 208.14/2
NPS-53-86-011

Added
P 208.113
MF5-55-86-00

NAVAL POSTGRADUATE SCHOOL
MONTEREY, CALIFORNIA 93943

R. C. AUSTIN
Rear Admiral, U. S. Navy
Superintendent

D. A. SCHRADY
Provost

Reproduction of all or part of this report is authorized.

This report was prepared by:

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER NPS-53-86-011	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle)		5. TYPE OF REPORT & PERIOD COVERED Technical Report April - August 1986
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Raul H. Mendez		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS Naval Postgraduate School Monterey, CA 93943		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Naval Postgraduate School Monterey, CA 93943		12. REPORT DATE September 1986
		13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS (If different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/ DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Vector Processor Floating Point Pipelines Scalar and Vector Performance Vector Ratio		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A summary of the scalar performance of the CRAY X-MP/4, Fujitsu VP-200 and NEC SX-2 is given. A description of the architecture, hardware and basic technology of SX-2.		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 63 IS OBSOLETE

S/N 0102- LF-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

THE PERFORMANCE OF THE NEC SX-2 SUPERCOMPUTER SYSTEM COMPARED
WITH THAT OF THE CRAY X-MP/4 AND FUJITSU VP-200.

Raul H. Mendez

Naval Postgraduate School, Monterey, California

Since the first delivery, late in 1983, of the Cray X-MP/2, Fujitsu VP-200 and Hitachi S-810/20 supercomputers, the race in high speed computers has considerably accelerated its pace. In 1984, both the Fujitsu VP-400 and the Cray X-MP/4 were first introduced and in the Fall of 1985 the Cray2 and the NEC SX-2 supercomputers were first brought into the market. The total number of installed systems including in-house systems number about 148 Cray systems, more than 40 CYBER CDC systems, about 44 VP systems and 13 Hitachi systems. So far, six NEC SX systems have been installed in Japan and one SX-2 system was delivered to the Houston Area Research Center this year, it is the first delivery of a Japanese system to an Academic Institution in the U.S. In this article we shall give an introduction to the SX-2 system, compare some of its features with those of the Fujitsu VP-200 (marketed in the USA by AMDAHL as the AMDAHL 1200) and CRAY X-MP/4 supercomputers (although not discussing in detail the latter systems) and survey some test data run on these three systems. The CRAY system will be referred as the X-MP or the X-MP/4, the Fujitsu-Amöahl machine will be referred to as the VP-200 or VP, and the NEC system as the SX-2 or SX.

It should be emphasized that our five benchmarks (fluid dynamics applications) codes are by no means detailed throughput tests and that our goal was not to obtain a detailed

performance profile but rather to sketch the salient features of the systems tested. Results on other benchmarks might yield different conclusions.

These results suggest that the SX-2 is a powerful processor of scalars and vectors, the fastest single processor in vector mode. In scalar mode the SX-2 was more than twice as fast as the VP-200 on all five benchmarks, and on the average about twice as fast as the X-MP/4 (these were all single processor tests and they were run on one single processor of the X-MP/4 that we tested). Before discussing in detail the three systems and results we shall review the importance of Amdahl's law in measuring the performance of a vector machine.

EFFECTIVE SPEED OF A VECTOR PROCESSOR

It has been widely recognized that the effective performance of a vector processor in real applications codes differ widely, often by an order of magnitude, from the advertised theoretical speed of the system. Gene Amdahl recognized the importance of scalar speed in estimating the total speed of the system. The time required to run the scalar (vector) portion of any give task or workload is inversely proportional to that system's scalar (vector) speeds. Since the total time required to run the workload is quite close to the net of these two times, it follows that no matter how fast the vector box of a supercomputer, the scalar portion will contribute to the total time. In real applications (medium vector ratios) the scalar contribution will dominate the total time. Therefore, unless the scalar speed is well balanced with the vector speed of a system, it can act as a

bottleneck to the system's performance (the dependence of total elapsed time on I/O processing speeds as well as OS overhead is analogous. Ours tests are all, however, CPU tests).

To illustrate the importance of scalar processing speed to the effective speed of a vector processor we shall use the above ideas to compare three hypothetical supercomputer systems, labelled A, B, C. In the following example the three systems are assumed to process a workload which is assumed to be 85% vector and 15% scalar. The scalars and vector speeds are assumed to be as listed in table 1, while the effective vector speeds entered in the last column are determined from Amdahl's law.

TABLE 1

----- Characteristics Speeds in MFLOPS of three hypothetical supercomputers for a workload which is 85% vector and 15% scalar -----			
System	Scalar Speed	Vector Speed	Effective Speed
A	2.5	300	15.9
B	5.0	150	28.1
C	10.0	300	56.2

The scalar speed of system B is assumed to be twice that of system A, while exactly the opposite relation holds between their vector speeds. As the table shows, despite the relatively high vector ratio (or vector rate) of this workload, in relative terms, the effective speeds of systems A and B more closely reflect

their scalar, rather than their vector speeds (the same can be said when comparing the effective speed of system C to that of systems A and B). This simple example points out that the effective speed of a supercomputer on a given application code is critically impacted by its scalar speed(A is an instance of a system with unbalanced scalar and vector speeds).

Consider now the effect on performance of compiler vectorizing capability. To illustrate the impact that different levels of compiler automatic vectorization has on performance assume that on the above workload the vector ratio yielded by system B can be increased to 90% a 5% gain over the vectorization yielded by the the other two compilers. Under this assumption the effective speed of system B becomes 38.5 Mflops. The speedup of system C over system B is thus reduced from 2 to 1.46. Thus the raw hardware power of system C can be partly balanced by the improved compiler sophistication of system B. Thus, a supercomputers system with a well balanced vector-scalar speed ratio is not effective unless it includes an adequate vectorizing compiler.

In addition to vector performance, compilers can significantly improve scalar performance. The CRAY CFT 1.15 compiler, for example, yields notable improvements in scalar performance over other versions of this compiler.

The above analysis has pointed out that the effective speed of a vector processor is influenced not only by the speed of its vector box but also by its scalar speed as well as by the sophistication of the system's compiler. We shall in particular

emphasize below the importance of compilers in our study of the performance of the SX-2, VP-200 and X-MP/4 supercomputers.

ARCHITECTURE AND HARDWARE OF THE SX-2 SYSTEM

This system design has targeted the scalar processing bottleneck and to implement that goal the SX designers have been guided by the ideas of distributed and RISC architectures (the number of vector instructions is 88 while that of scalar instructions is 83).

The system consists of two processors that can operate concurrently, the control and arithmetic processors. The control processor runs the operating system, the compiler and executes other supervising tasks. The control processor's design is based on that of NEC's ACOS mainframe computer, a general purpose computer with an advertised performance in the 30 MIPS range, for the single processor configuration.

The arithmetic processor of the SX-2 consists of two subunits each running at a clock speed of 6 nsec. The scalar unit includes a set of four fully segmented pipelines including floating point add and multiply. Instruction processing is accelerated by a 2k byte instruction buffer and scalar operands memory accesses are speeded up by a 64 K-byte cache , as in the VP-200 system (a single processor of the X-MP/4 uses its 64 T registers to store intermediate results). Scalar operands are directed from the general purpose cache to the scalar registers (128 of these are available, there eight scalar S registers in one processor of the X-MP/4) and from there routed to the scalar

pipelines. The SX as the X-MP processes scalars, in pipeline fashion, and this feature as well as the large number of scalar registers should have a direct impact on scalar performance.

The vector unit consists of four sets of vector pipelines, netting a total of eight floating pipes (four add and four multiply). Vector transfer rates are speeded up by a set of forty vector registers, each with a capacity of 256 elements, for a total capacity of 80k bytes (as opposed to 64k bytes on the VP-200 and 8k in one processor of the X-MP/4).

The computing rate is sustained by eight load and four store pipes which cannot operate concurrently (all load and store pipes are 64 bits wide). When chaining is possible the maximum vector computing rate is in principle eight results every clock (every 6 nsec) , as opposed to four results every 7 nsec in the VP-200 and two results every 9.5 nsec in one processor of the X-MP/4. A masking pipeline is available for the implementation of conditional vector operations. As in the X-MP/4 and VP systems special purpose hardware is used in gather scatters operations.

MEMORY

The SX-2's memory has a maximum capacity of 256 megabytes, the same maximum capacity as in the VP-200 while the maximum is 128 Megabytes on the X-MP/4. The degree of interleaving is 64 banks in the X-MP/4 and effectively 256 on the SX-2, the same level of interleaving as in the VP-200. In addition to the main memory, the control processor of the SX-2 includes 64 Megabytes of local

memory (both local and main memory are addressable by the control processor).

The bandwidth of the main memory as stated earlier is 8 words per clock or 1.33 gigawords as opposed to 315 million words on one processor of the X-MP/4 (three words per cycle) and 565 million words per second on the VP-200. On the other hand, a load operation, that is a fetch from memory to vector registers, requires 36 clocks (216 nsec) as opposed to 14 clocks (133 nsec) in the X-MP. Longer startup times are needed for vector operations and thus the vector performance of the X-MP/4 on short lengths should be superior to that of the SX-2.

The main memory is supported as in the X-MP by an SSD device (no SSD is available on the VP-200). The maximum capacity of the SSD is 2 gigabytes and 1 gigabyte on the X-MP/4. The transfer rate between the main memory and the SSD is 1.3 gigabytes per second in the SX-2 and 2 gigabytes per second in the X-MP/4. The availability of the SSD should have considerable impact on I/O handling but none of our tests tested this capability.

EFFECTIVE VECTOR PERFORMANCE

The vector performance of a supercomputer is determined not only by the rate at which operands can be processed by the pipes within the vector box but also by the flow rate of these operands between memory and pipes. Thus, as scalar speed can slow down the effective speed of a vector processor, slow memory accesses can become a major bottleneck in vector performance. Memory reads and writes can proceed in three different modes on a

vector processor. Contiguous, strides and gather-scatters. The first two accesses refer to accessing equispaced memory locations (spaced by one word in the contiguous case) while the last refers to memory accesses governed by a list vector, which accesses memory locations in an irregular manner. The mix of these three types of accesses on a given workload as well as the ratio of operations to accesses determine the effective vector speed (in general gather-scatter accesses are the slowest and contiguous are the fastest).

Our benchmark data well as performance data from simple vector operations and kernels published elsewhere lead to the following observations. All three systems handle contiguous accesses at their maximum bandwidth rate. Equispaced memory access with even stride slow down considerably on both the Fujitsu and NEC systems, while the Cray handles most stride memory accesses at full bandwidth speed. On the SX-2, the slow-down depends not only on the stride but also on the ratio of vector operations to memory accesses within a given vector loop (odd strides accesses were not tested). Memory strides which are powers of two, as those needed in FFT routines processing a number of data points which is also a power of 2 slow down considerably on the SX-2. The advantage of the Cray system in regards to equispaced memory accesses results from the fast cycle time of its memory. In one processor of the X-MP/4 four clocks (38 nsec) must elapse between memory accesses to the same bank, while 13 clocks (78 nsec) are needed in the SX-2. Thus, a memory fetch to the same bank can result in a longer wait in the NEC system. The number of

banks is however four times that of the X-MP/4 system. The X-MP/4's faster memory cycle times results directly from its use of ECL bipolar RAMs in main memory as opposed to the MOS static RAMs used in the NEC system. The three systems include the necessary hardware to handle gather-scatter memory accesses, however, but we have not tested this type of memory access.

BASIC TECHNOLOGY USED IN THE SX-2 SYSTEM

The achievement of the 6 nsec clock in the SX-2 is possible through the implementation of very fast densely packaged logic. Liquid convection technology allows high gate density packaging.

The main memory devices are 64 Kbit static RAMs with 40 nsec access times, while 256 dynamic RAMs with 120 nsec access times are used in the SSD. Vector registers and cache are implemented in 1 Kbit 3.5 nsec access time bipolar LSI. Logic is implemented in 1000 gate arrays chips with gate delays of 250 picoseconds.

Memory is packaged in 3-d modules, each with a capacity of two megabytes. Logic is cased in special purpose thermal cooling modules which house up to 36 LSI, for a maximum 36000 gates per package. Air cooling is used to cool the main memory device and a water cooling convection system is used to convect the over 200 Watts dissipated by each LSI package (there are in total 92 of these packages).

PERFORMANCE

Five fluid dynamics applications codes gathered from different sources were used as testing instruments. The same five programs

were used in an earlier comparison study of the Fujitsu VP-200 and Cray X-MP systems. These codes do not represent any given workload and are characteristic only of the types of fluid dynamics modeling used in these programs. Two of them MHD-2D and SHEAR3 have been used extensively in turbulence simulations in two and three dimensions and developed on Cray systems. BARO is a two dimensional shallow water mode of the atmosphere, which has been developed on the CDC CYBER 205. EULER is a one-dimensional spectral code used to model the shock-tube problem, developed on a TI's ASC system and VORTEX is a particle simulation code developed on an IBM 3033 main-frame.

In our timings the following ground rules were used. Codes BARO and VORTEX were run unmodified in all three systems, slight tuning was allowed in EULER (up to twenty lines) and about the same finite amount of time was given to the three makers to tune the other two codes, MHD-2d and SHEAR3.

Compilers used in our testing are as follows. The SX-2 vector timings were obtained with versions 20 and 24 of the compiler, the vector results with the latter version are faster and thus our discussion of vector performance will be based on these timings. Scalar timings analysis is based on data obtained with version 20 of the compiler (versions 20 and 24 yield nearly the same scalar performance. Similarly, versions V10L10 and V10L20 of the VP-200 were used in vector mode, but analysis of the results on this mode are based on the V10L20 compiler. Because the most recent version of the compiler V10L31 yields notable improvements in scalar (and nearly the same performance in vector

mode) this version was used in our analysis of scalar performance of the VP-200. The vector and scalar timings of the X-MP/4 were obtained with version CFT1.15 of the CRAY compiler. All runs were obtained in dedicated mode, at the NEC Fuchu plant in Japan, the Sunnyvale AMDAHL facility in California and the Mendotta Heights CRAY facility in Minnesota.

SCALAR PERFORMANCE

One of the strongest features of the SX system lies in its strong scalar processing power. Table 2 shows that the floating point operations run faster on the SX-2 than on the other two systems. However, the speed up obtained in our tests is far from that suggested by these speeds alone. In fact the fast scalar performance of the SX-2 system is the result not only of the fast clock but of other features such as the large number of scalar registers, pipelined functional units and the ability of the compiler to schedule scalar operations with a high degree of concurrency. The scalar unit's cache memory, also available on the VP-200, is also an important performance factor. The impact of the faster SX-2 clock is felt on transfers of data from memory when a cache miss takes place (the VP-200 scalar clock is 14 nsec versus 6nsec on the SX-2).

TABLE 2
TIMINGS OF FLOATING POINT OPERATIONS

	SX-2	VP-200	X-MP
	lclock=6nsec	lclock=14nsec	lclock=9.5nsec
Operation	nsec (clocks)	nsec(clocks)	nsec(clocks)
Floating Point Add	36 (6)	42 (3)	57 (6)
Floating Point Multiply	54 (9)	56 (4)	66.5(7)

RESULTS IN SCALAR MODE

RESULTS IN SCALAR MODE

In two of the codes, SHEAR3 and EULER, the SX-2 was about 2.6 times faster than one processor of the X-MP/4. Most of the work in these two codes is done on FFT routines, processing arrays that can be kept in cache on the SX-2 and VP-200 throughout the computation. The VP-200 processes these two codes faster than one processor of the X-MP/4 but it is slower than the SX-2 by a factor of 2.21 in EULER and 2.50 in SHEAR3 (this last result was obtained using the V10120 compiler).

In MHD-2D most of the work is done on an FFT routine processing two-dimensional 256x256 arrays which cannot be kept in cache. Memory conflicts, since the strides are powers of two, slow down the SX-2 and VP-200 vis a vis the X-MP/4. In this program one processor of the X-MP/4 and the SX-2 yielded identical times, while the SX-2 was 2.04 times faster than the VP-200.

As in MHD-2d, in BARO most of the work is done on arrays too large to be kept cache. The memory accesses also slow down large to be kept in cache. The memory accesses also slow down its performance on the VP-200 (this program suffered a performance degradation when run on a VP-100 with half the number of banks used in the VP-200). The SX-2's speedup over one processor of the X-MP/4 is 1.79 and it is 2.28 times faster than the VP-200 on this code.

In VORTEX the speedup of the SX-2 over one processor of the X-MP/4 is 1.80 and the SX-2's speedup over the VP-200 is 2.01. Performance analysis in this code is more complex than in the other benchmarks

TABLE 3

SCALAR TIMINGS IN SECONDS

V/S stands for VP-200 to SX-2 timing ratio, and X/S stands for X-MP/4 to SX-2 timing ratio

Code	SX-2 vers.20	VP-200 V10L31	X-MP/4 CFT1.15	V/S	X/S
BARO	398.8	910.7	713.7	2.28	1.79
EULER	2.9	6.4	7.5	2.21	2.59
MHD2-D	18.4	37.5	18.4	2.04	1.00
SHEAR3	65.7	164.4	172.2	2.50	2.62
VORTEX	76.7	154.4	138.2	2.01	1.80

VECTOR PERFORMANCE

As described above the scalar speed of a vector processor plays

an important role in its overall performance unless the vector ratio of the workload is close to 100%. In performance studies of supercomputers computing the vector speed of a given benchmark in each system accurately is generally difficult. Data on the SX's ANALYZER SUMMARY of each code facilitates estimating vector and scalar speeds on the SX-2, in particular the vector operation ratio given as output by the ANALYZER, can be used to estimate the vectorization ratio in each code. Three of our test programs, BARO, MHD-2d and VORTEX were highly vectorized by the three systems' compilers, the other yielded medium vector ratios in all three systems.

We shall see below that our benchmark data provides an indirect assessment of the performance of the three systems in the range from short to moderately long vectors as well as with medium to high vector ratios. Performance with contiguous and strided accesses also were indirectly tested by our benchmarks. In regard to the latter it should be clarified that three of the codes ran a significant part of the work on FFT routines and that the two types of FFT's used (the same FFT routine was used in MHD-2d and SHEAR3 and a less efficient version was used in EULER) have not been specially coded to vectorize. In fact, the FFT used in the program EULER, includes the type memory access (strides which are powers of 2) which most adversely affect vector speed because of the resulting bank contention. We have opted for not using the systems' FFT libraries because our objective was not to test specific aspects of the systems (such as Library FFTs) but rather to test their ability to process more or

less typical FORTRAN codes.

COMPILER PERFORMANCE

Table 4 shows the results of running the five benchmark codes in vector mode on the three different systems. The benchmark set has been run on each system under two different versions of the compiler on the indicated dates. Timings improvement with each compiler version were strictly due to the compilers, no code changes were allowed in the benchmark set between the two timings.

TABLE 4
TIMINGS IN VECTOR MODE USING TWO DIFFERENT COMPILERS

CODE	SX-2		VP-200		X-MP	
	Ver.20	Ver.24	V10L10	V10L20	CFT1.13	CFT1.15
	11/85	4/86	1/86	1/86	2/84	2/86
	(sec)		(sec)		(sec)	
BARO	19.4	19.6	38.2	38.2	76.3	70.5
EULER	1.9	2.0	5.3	4.6	3.1	2.9
MHD-2D	1.6	1.2	2.0	2.0	4.3	3.7
SHEAR3	44.5	40.0	72.1	71.6	72.7	58.1
VORTEX	7.2	6.1	13.7	12.4	NA	13.9

The compilers performance on our benchmarks suggest that the level of the three systems compilers may be roughly comparable. The VP-200 and SX-2 version 24 compilers include nearly the same

automatic vectorization features, with the CFT 1.15 not far behind. The main feature of the VP compiler not yet available on the SX-2 compiler is the vectorization of some types of nested double loops.

In program BARO the V10120 compiler vectorized 66 loops, the CFT1.15 61 loops and the version 24 of the SX-2 compiler, 62 loops (the advantage of the VP compiler was due in this case to four double loops). A similar situation occurs in VORTEX, the VP vectorized 25 loops the SX 23 and the X-MP 23 loops. In code Euler the VP compiler vectorized one more loop than the SX-2's, fifty-one versus fifty. The non-vectorized loop with length 4, a length below the break-even-point between scalar and vector on the SX-2, defaulted to scalar mode. The CFT1.15 vectorized, after hand restructuring, the same fifty one loops vectorized by the VP compiler, because of loop splitting these fifty-one loops were turned into fifty five loops. In SHEAR3 after some restructuring 38 loops were vectorized on the VP, 36 on the SX and the X-MP vectorized 35 loops. In MHD-2D after restructuring 28 loops were vectorized by the VP200, 28 by the SX-2 and 26 by the CFT.

RESULTS IN VECTOR MODE

We summarize in table 5 characteristics speeds of the codes tested. Next a summary of the performance on each of the VP and X-MP systems vis a vis the SX-2 is given in table 6. The data on these tables is surveyed first and then each code's data is discussed in some detail.

The vector ratio on each system can be estimated by considering

the ratio of performance in scalar and in vector mode. Thus, from table 5 we can infer that the codes with the highest vectorization ratios are BARO, VORTEX and MHD-2D. These speedups slow down considerably on codes EULER and SHEAR3.

TABLE 5

RATIO OF SCALAR TO VECTOR TIMINGS ON EACH CODE

	BARO	VORTEX	EULER	MHD-2D	SHEAR3
SX-2	20.3	12.4	1.6	15.3	1.6
VP-200	29.0	11.8	1.4	21.7	2.3
X-MP	10.9	9.9	3.10	10.6	3.3

Table 6 summarizes the relative speed up of the SX-2 relative to the other two systems in vector mode (combined scalar and vector performance). Notice that the relative speedup of the VP-200 vis a vis the SX-2 is with one exception (EULER), quite consistent ranging from 1.7 to 2.0. There is a wider performance range in the performance of one processor of the X-MP/4 relative to that of the SX-2, from 1.5 to 3.6.

TABLE 6
RELATIVE SPEEDUP OF THE SX-2 OVER THE VP-200 AND X-MP
IN VECTOR MODE

	BARO	VORTEX	EULER	MHD-2D	SHEAR3
VP-200	1.9	2.0	2.3	1.7	1.8
X-MP	3.6	2.3	1.5	3.1	1.5

We proceed to discuss these results beginning with the code with the highest effective to scalar performance ratio.

BARO

The sixty-one loops of this code vectorized in all three systems amount to more than 99% of the total work. Memory accesses are contiguous and vector length moderately long at 300. Table 6 shows that in this program the speed of the SX-2 is 1.9 times that of the VP-200 and 3.6 times that of one processor of the X-MP/4. These ratios are not far from the ratio's in maximum vector throughput of these systems. It is noteworthy also that the VP-200 is the system with the highest vector/ scalar speed ratio, the VP-200 executes this code in vector mode twenty nine times as fast as in scalar mode. These speedups are about 11 and 20 on the X-MP/4 and SX-2). In program BARO the effective speed up of the SX-2 over the VP-200 is 1.94 while the scalar speedup is 2.78. The effective speedup is close to the ratio of vector throughputs. Performance is dominated by vector speeds and the scalar advantage of the SX-2 does not play a role.

VORTEX

The code VORTEX is a particle code which simulates the dynamics of a 1-D Vortex sheet by means of discrete vortices. In VORTEX as in BARO, memory accesses are contiguous and the vector ratio is quite high (99.% vector operation ratio, according to the SX Analyzer). Indeed, in VORTEX as in BARO, the compiler performance of the three systems is nearly the same and though the VP

compiler vectorized two more loops than the SX these loops amounted to less than 1% of the total CPU time on the SX-2. Unlike BARO, the vector lengths in the two most CPU bound loops of VORTEX increase from 20 to 500 in strides of 1. Due to the strength of the X-MP/4 in handling short vectors, despite the high vector ratio the performance of the VP-200 and the X-MP/4 are close at 12.4 and 13.9 sec respectively. The SX-2's timing is in this case 2.02 times faster than the VP-200 and 2.28 times faster than the X-MP. Thus, although a high degree of vectorization is obtained on this code by the three systems, the short vector lengths slow down the SX-2 and VP-200. Thus, relative to these two systems, the X-MP/4 performs better in VORTEX than in BARO (both with vector ratios of nearly 99% in the three systems).

FFT CODES

The remaining three codes spent a significant part of the total CPU work in FFT routines. As was mentioned above, the performance of the three systems on these three codes should not be interpreted as representative of their performance in handling FFT work.

In vector mode on the SX-2, FFT work amounts to 69%, 57% and 31% on EULER, MHD-2d and SHEAR3 respectively (these rates are not estimates but are derived by the Analyzer from actual timings). In code Euler, memory conflicts slow down the speed of the SX-2 in vector mode to nearly 2/3 of its scalar speed while processing the FFT routine (1.1 sec to 1.5 sec). As mentioned before this performance degradation is the result of the adverse powers of 2

strides used in Euler's FFT routine. Memory conflicts have an effect also on the SX-2 MHD-2d and SHEAR3 performance, however their impact on vector speed is less drastic than in EULER's case (different FFTs are used in Euler than in SHEAR3 and MHD-2d). The longer vector lengths used in MHD-2d (typical vector length is 256) conceal the impact of the strides on the SX's performance in vector mode. In MHD-2d, the FFT routine in vector mode runs 22.1 times faster than in scalar mode. The effect of the strides is particularly apparent when the vector length is short as in SHEAR3 (typical vector length is 16). In this test the SX-2 in vector mode processes the same FFT routine used in MHD-2d 2.5 faster than in scalar mode.

EULER

Because of the type of FFT used in this code and because it is a one-dimensional code this benchmark is perhaps, within the benchmark set, least representative of the codes used in large scale computing. Despite the fact that up to twenty lines of FORTRAN tuning was allowed, the resulting code is virtually the same on all three systems, tuning was restricted to compiler directives and restructuring of the same loops. The same fifty loops were vectorized by the three compilers and we shall assume that the Euler's vector ratio is nearly the same in all three systems. Euler's vector operation ratio is 73% on the SX-2. In vector mode on this code the SX-2 was 2.30 times faster than the VP-200 and 1.45 times faster than the X-MP. On this code the ratio of timings in scalar to vector mode is 1.37 on the VP-200 and 1.55 on the SX-2 and 3.10 on one processor of the X-MP/4.

Thus, the X-MP/4 is the least affected by the power of two stride memory accesses and the VP-200 the most. It is noteworthy that the SX-2 in scalar mode at 2.9 sec, outperformed the VP-200's timing in vector mode, 4.6 sec, and matched the timing in vector mode of one processor of the X-MP/4 at 2.9 sec.

MHD-2d and SHEAR3

The codes MHD-2d and SHEAR3 are two and three dimensional turbulence fluid dynamics simulation based on spectral techniques. Thus, again the FFT routine (differently coded) is the most active in CPU usage. On both these codes limited tuning was permitted on the three systems tested and the vector ratios in the three systems may not be the same.

According to the SX-2's ANALYZER the vector operation ratio on MHD-2d is 99%. Typical vector length in this code is 256. In this code the SX-2 is 1.67 times faster than the VP-200 and 3.08 times faster than the X-MP/4. The longer vector lengths in this program as well as the high vector ratio allow effective use of the vector pipes on both the VP-200 and SX-2 systems and their vector speeds are only partly reduced by the strides. The ratio of effective speed to scalar speeds is 21.7 times on the VP-200, 15.2 on the SX-2 and 10.6 on one processor of the X-MP/4.

SHEAR3 is a 3-D calculation using the same FFT routine used in MHD-2D. The vector operation ratio according to the SX-2 ANALYZER is 89% on this code. The SX-2 is 1.45 times faster than one processor of the X-MP/4 and 1.79 times faster than the VP-200. In this case the strong performance of the X-MP/4 with

short vector becomes apparent as does the slow down of the VP and SX-2 systems when handling even strides and short vector loops. In this code the ratio of effective to scalar performance on the SX-2 is 1.64, 2.30 on the VP-200 and 3.28 on the X-MP. It is noteworthy that the scalar performance of the SX-2 at 65.8 sec is in this case faster than the vector performance of the VP system's 71.6 sec in vector mode.

SUMMARY OF RESULTS IN VECTOR MODE

The speedup of the SX over the VP-200 is with exception of program Euler (2.3 speedup) between 1.7 and 2.0. In EULER, memory conflicts slow down the VP-200 to 1.37 of its scalar performance. The speedup of the SX-2 over one processor of the X-MP/4 is less consistent, varying from 1.45 to 3.60. The highest speedups 3.60 and 3.08 are associated with the high vector ratios and vector lengths present in programs BARO and MHD2d. In Vortex although the vector ratio is high the calculation includes short vectors and the speedup is reduced to 2.28. This ratio is reduced further as the vector length is shortened and the memory accesses are the even powers strides found in Euler. The lowest value of this speedup, 1.45, occurs with the program SHEAR3, in this case the calculation involves short vector and even strides.

CONCLUSIONS

1) The SX-2 system is an outstanding system in regard to the processing of scalars. The SX-2 was in scalar mode, about twice as fast as one processor of the X-MP/4 and more than twice as

fast as the VP-200.

2) In vector mode the SX-2 was up 3.6 times faster than a single processor of the X-MP/4 for a vector length of 300 as well as vector ratio of 99%. For short vector lengths(16) and even strides the SX-2 was 1.5 times faster.

3) The SX-2's speed up in vector mode over the VP-200 was between 1.7 and 2.0 with one exception (2.30).

4) The compiler performance of the SX-2 (version 24) is quite close to that of VP's V10L20 and the CFT1.15 is not far behind these two compilers in vectorization capability.

5) The X-MP/4 system is the least affected by short vectors and by even strides.

6) I/O and O/S overhead have not been accounted for. A performance study including the latter two components in the total performance of the systems may lead to different results.

REFERENCES

1) Bucher, I. Y., "The Computational Speed of Supercomputers", Proc. ACM/SIGMETRICS Conf. on Measurements and Modelling of Computer Systems, Aug. 1983, pp. 151-165.

2) Lubeck, O., Moore, J., Mendez R., "A Benchmark Comparison of Three Supercomputers: Fujitsu VP-200, Hitachi S-810/20 and Cray X-MP/2. Proc. First International Conference on Supercomputing Systems, Dec. 1985.

DISTRIBUTION LIST

	<u>No. of Copies</u>
Defense Technical Information Center Cameron Station Alexandria, Virginia 22314	2
Superintendent ATTN: Library, Code 0142 Naval Postgraduate School Monterey, CA 93943	2
Research Administration Code 012 Naval Postgraduate School Monterey, CA 93943	1
Chairman, Dept. of Mathematics Code 53 Naval Postgraduate School Monterey, CA 93943	1
Assistant Professor Raul Mendez Code 53Mu Deptment of Mathematics Naval Postgraduate School Monterey, CA 93943	10

DUDLEY KNOX LIBRARY



3 2768 00342573 7